

**APPLICATION**  
**FOR**  
**UNITED STATES LETTERS PATENT**

**TITLE:**                    **SYNCHRONIZATION OF DATABASES USING FILTERS**  
**APPLICANTS:**        **DAVID J. BOOTHBY AND DAVID W. MORGAN**

"EXPRESS MAIL" Mailing Label Number EM53918475C5

Date of Deposit March 5, 1968

I hereby certify under 37 CFR 1.10 that this correspondence is being deposited with the United States Postal Service as "Express Mail Post Office To Addressee" with sufficient postage on the date indicated above and is addressed to the Assistant Commissioner for Patents, Washington, D.C. 20231.

Lisa C. Gray  
Lisa C. Gray

SYNCHRONIZATION OF DATABASES USING FILTERS

Cross Reference To Related Applications

5           This application is a continuation in part of  
"Synchronization of Databases with Date Range," Serial no.  
08/748,645, filed November 13, 1996 .

Reference to Microfiche Appendix

10           An appendix (appearing now in paper format to be  
replaced later in microfiche format) forms part of this  
application. The appendix, which includes a source code  
listing relating to an embodiment of the invention, includes  
\_\_\_\_\_ frames on \_\_ microfiche.

15           This patent document (including the microfiche  
appendix) contains material that is subject to copyright  
protection. The copyright owner has no objection to the  
facsimile reproduction by anyone of the patent document as  
it appears in the Patent and Trademark Office file or  
records, but otherwise reserves all copyright rights  
20           whatsoever.

Background

          This invention relates to synchronizing databases.

25           Databases are collections of data entries which are  
organized, stored, and manipulated in a manner specified by  
applications known as database managers (hereinafter also  
referred to as "Applications"; hereinafter, the term  
"database" also refers to a database manager combined with a  
database proper). The manner in which database entries are  
organized in a database is known as the data structure of  
30           the database. There are generally two types of database  
managers. First are general purpose database managers in  
which the user determines (usually at the outset, but  
subject to future revisions) what the data structure is.

These Applications often have their own programming language and provide great flexibility to the user. Second are special purpose database managers that are specifically designed to create and manage a database having a preset data structure. Examples of these special purpose database managers are various scheduling, diary, and contact manager applications for desktop and handheld computers. Database managers organize the information in a database into records, with each record made up of fields. Fields and records of a database may have many different characteristics depending on the database manager's purpose and utility.

Databases can be said to be incompatible with one another when the data structure of one is not the same as the data structure of another, even though some of the content of the records is substantially the same. For example, one database may store names and addresses in the following fields: FIRST\_NAME, LAST\_NAME, and ADDRESS. Another database may, however, store the same information with the following structure: NAME, STREET\_NO., STREET\_NAME, CITY\_STATE, and ZIP. Although the content of the records is intended to contain the same kind of information, the organization of that information is completely different.

Often users of incompatible databases want to be able to synchronize them with one another. For example, in the context of scheduling and contact manager Applications, a person might use one application on a desktop computer at work while another on his handheld computer or his laptop computer while away from work. It is desirable for many of these users to be able to synchronize the entries on one with entries on another. U.S. patents of the assignee hereof, Puma Technology, Inc. of San Jose, California (U.S.

Patent No. 5,392,390, hereinafter, "the '390 patent", incorporated by reference herein; and U.S. Patent No. 5,684,990, filed on January 11, 1995, incorporated by reference herein) show two methods for synchronizing incompatible databases and solving some of the problems arising from incompatibility of databases.

#### Summary

In one general aspect, the invention features a computer program for synchronizing at least a first and a second database. A plurality of records of the first database fitting a selected criterion are identified. At least one of the identified records of the first database is then synchronized with a record of the second database.

In another general aspect, the invention features a computer program for synchronizing at least a first and a second database. On a computer display, a record selection criteria input region is displayed for a user to input a record selection criteria. Then, the first database is synchronized with the second database using the record selection criteria.

Preferred embodiments may include one or more of the following features.

Records representative of the records of the first and second databases during a prior synchronization are stored in a history file. In that case, when synchronizing the identified records of the first database with the records of the second database, the history file is used.

Records of the second database may also be identified based on a selected criterion. In that case, the identified records of the first database are synchronized with the identified records the second database.

Records of the first and second databases may include text, number, boolean, binary, date, and time fields. The criteria for identifying the records in turn may include text, number, boolean, binary, date and time  
5 criteria with which the fields of the records and databases are compared.

The first database can be located on a first computer and the second database located on a second computer. At the first computer, it is then determined  
10 whether a record of the first database has been changed or added since a previous synchronization, using a first history file located on the first computer including records representative of records of the first database at the completion of the previous synchronization. If the record  
15 of the first database has not been changed or added since the previous synchronization, information which the second computer uses to identify the record of the first database to be unchanged is sent from the first computer to the second computer. Additionally, the identification of the  
20 records of the first database based on the selected criterion may be performed at either the first or second computer.

Based on data reflecting whether the records of the first database have been added or changed since a previous  
25 synchronization, it may be determined whether the records of the first database have been changed or added since a previous synchronization. If one of the records of the first database has not been changed or added since the previous synchronization, a synchronization with records of  
30 the second database using a record representative of the one record at the time of a previous synchronization is performed. The representative record is stored in a history

file containing records reflecting the contents of records of the databases at the time of a previous synchronization.

A second plurality of the records of the first database failing to fit the selected criterion may be  
5 deleted.

The selected criterion may have a current value during a current synchronization being different from a previous value during a previous synchronization. In that case, a plurality of records of the second database may be  
10 updated, based on results of the synchronization, where the plurality of records of the second database fit the previous value of the selected criterion but fail to fit the current value of the selected criterion.

A third database may be synchronized with the second  
15 database by identifying a plurality of records of a third database fitting a second selected criterion and synchronizing at least one of the identified records of the third database with a record of the second database. The record of the second database can include a code identifying  
20 the record as having originated from the third database.

A record selection criteria may be transmitted to a database manager which manages the first database and the database manager may select records of the first database fitting the record selection criteria. The database manager  
25 may then transmit the selected records to the synchronization program. The records of the first database fitting the record selection criteria may also be selected at a synchronization program.

The selected criterion may be, for example, a filter  
30 or filter expression which a record must match or fit in order for that record to pass the filter expression.

Embodiments of the invention may include one or more of the following advantages.

Users of various embodiments of the invention can use those embodiments to achieve a variety of ends. For example, handheld computers typically have limited storage capacity. Using the filtering capability of some  
5   embodiments, a user can limit the records stored on the handheld computer to only those records which fit a selected filter.

10   A user can also use a filter in some embodiments to increase the speed of synchronization. For example, it may be the case that the data transfer link between the two computers has a low data transfer rate. Therefore, the user by using a filter reduces the number of records to be transferred from one computer to the other.

15   A user can also use the filtering mechanism to synchronize some of the records of his or her database with the related records of another user's database, without affecting the unrelated records. For example, consider the situation where two database users work on a shared project or take a joint business trip on behalf of their enterprise.  
20   These two database users may desire to synchronize their databases but only with respect to those records relating to that project or trip. By using a filter, they may limit the synchronization between the two databases to records which relate to the project or the trip, without affecting other  
25   records in their respective databases.

The invention may be implemented in hardware or software, or a combination of both. Preferably, the technique is implemented in computer programs executing on programmable computers that each include a processor, a  
30   storage medium readable by the processor (including volatile and non-volatile memory and/or storage elements), at least one input device, and at least one output device. Program code is applied to data entered using the input device to

perform the functions described above and to generate output information. The output information is applied to one or more output devices.

Each program is preferably implemented in a high level procedural or object oriented programming language to communicate with a computer system. However, the programs can be implemented in assembly or machine language, if desired. In any case, the language may be a compiled or interpreted language.

Each such computer program is preferably stored on a storage medium or device (e.g., ROM or magnetic diskette) that is readable by a general or special purpose programmable computer for configuring and operating the computer when the storage medium or device is read by the computer to perform the procedures described in this document. The system may also be considered to be implemented as a computer-readable storage medium, configured with a computer program, where the storage medium so configured causes a computer to operate in a specific and predefined manner.

Other features and advantages of the invention will become apparent from the following description of preferred embodiments, including the drawings, and from the claims.

#### Brief Description of the Drawing

Figure 1 shows two computers connected via a data transfer link.

Figure 2 is a schematic drawing of the various modules constituting an embodiment of a synchronization program.

Figure 3 is a representation of a workspace data array used by the synchronization program of Figure 2.



Figure 4 is pseudocode for a Translation Engine Control Module of the synchronization program of Fig. 2.

Figure 5 is pseudo code for the steps taken by Parameter Table Generator module of the synchronization program of Fig. 2.

Figure 6 shows a filter selection graphical user interface (GUI) window.

Figure 7 shows a filter name input graphical user interface (GUI) window.

Figures 8 and 9 show a filter criteria input graphical user interface (GUI) window.

Figure 10 shows a table detailing semantics of a filter language used in the synchronization program of Figure 2.

Figure 11 is pseudocode for loading a history file.

Figure 12 is pseudocode for the steps taken by a translator to load records of a remote database.

Figure 13 is pseudocode for the steps taken by the synchronizer module of the synchronization program of Figure 2 for performing Conflict Analysis and Resolution when synchronizing using a filter expression.

Figure 14 is pseudocode for the steps taken by a translator to unload records to a remote database.

Figure 15 is pseudocode for the steps taken by a translator to unload records to a local database.

Figure 16 is a schematic drawing of the various modules constituting an embodiment of a distributed synchronization program.

#### Description

We will describe embodiments of the invention in detail below, but briefly, referring to Figs. 1 and 2, a synchronization program 100 runs on a local computer 20

(e.g. a desktop or server computer) which is typically connected to a remote computer 22 (e.g. a handheld or notebook computer) via a data transfer link 24 enabling the computers to transfer data between them. Data transfer link  
5 24 may be a serial infrared link, serial cable, modem and telephone line combination, or other such data transfer links. Each of the local and remote computers stores a corresponding local or remote database, which may, for example, be a scheduling database (such as those sold under  
10 the tradenames Microsoft Schedule+ and Lotus Organizer).

Synchronization program 100 synchronizes the records of the local and remote databases typically using a history file that contains records reflecting the records of the two databases at the end of a previous synchronization. The  
15 synchronization program uses the history file to determine, for example, which records have been changed, added or deleted since the previous synchronization and which records of the two databases correspond to one another.

Synchronization program 100 allows the user to input  
20 a filter expression. Generally, a filter or filter expression may considered to be a set of conditions or criteria which a record must match or fit in order for that record to pass the filter expression. A record that does not fit those criteria therefore fails the filter.

25 Synchronization program 100 uses the filter expression to identify and mark which records of the local and remote databases pass or fail the filter expression. The two databases are then synchronized. After synchronization, synchronization program 100 uses the  
30 results of synchronizing the two databases to add, modify, or delete the records of the two databases. At this point, the user can, for example, select to have those records falling outside the filtering criteria to be deleted, not to

be affected at all by the synchronization program, or be treated in other manner, as will be described below. In other embodiments, synchronization program 100 uses the filter expression to identify and mark the records of only one of the databases.

We will now describe in detail the structure of synchronization program 100 and the method it uses to synchronize the local and remote databases using filter expressions. Fig. 2 shows the relationship between the various modules of an embodiment of synchronization program 100. Translation Engine 1 comprises Control Module 2 and Parameter Table Generator 3. Control Module 2 is responsible for controlling the synchronizing process by instructing various modules to perform specific tasks on the records of the two databases being synchronized. (Fig. 4 shows the steps taken by this module.)

Parameter Table Generator 3 is responsible for creating a Parameter\_Table 4 which is used by all other modules for synchronizing the databases. Generally, Parameter\_Table 4 stores various information which may be used by the modules of the synchronization program. The information stored in Parameter\_Table 4 includes user preferences, the names and locations of the databases, and the names and locations of various files stored on disk including the name and location of the history file from the previous synchronization. Parameter Table Generator 3 also provides the user with various graphical user interface windows for inputting filter expressions to be used during synchronization. Parameter Table Generator 3 converts such user input filter expressions or previously stored filter expressions into data structures which may then be used by various modules of synchronization program 100 during synchronization. The steps taken by Parameter Table

Generator 3 in relation to filter expressions will be described in detail below.

Synchronizer 15 has the primary responsibility for carrying out the core synchronizing functions. It is a table-driven code which is capable of synchronizing various types of databases whose characteristics are provided in Parameter\_Table 4. Synchronizer 15 creates and uses workspace 16 (also shown in Fig. 3), which is a temporary data array used during the synchronization process.

Synchronization program 100 has two translator modules 5 and 9 which are generally responsible for data communication between synchronization program 100 and databases 13 and 14. Translator (L\_translator) 5 is assigned to the local database (L\_database) 13 and translator 9 (R\_translator) to the remote database (R\_database) 14. Each of the database translators 5 and 9 comprises three modules: reader modules 6 and 10 (L\_reader and R\_reader) which load (or read) records from databases 13 and 14; unloader modules 8 and 12 (L\_unloader and R\_unloader) which analyze and unload records from workspace 16 into databases 13 and 14; and sanitizing modules 7 and 11 (L\_sanitizer and R\_sanitizer) which analyze the records of the opposing database when they are loaded into the workspace and modify them according to rules of data value of the modules's own database. Briefly stated, rules of data value are generally rules that define the permitted content of the fields of the records of a database. An example of such a rule would be that no more than 100 characters may be present in a field, or that content of a field designating a priority for a "to do" item should be limited to 1, 2, or 3. Sanitizing a record is to change the content of the fields of a record of one database to conform to the rules of data value of another database. Rules of

data value and sanitization are described in detail in the following commonly owned U.S. patent applications, incorporated in their entirety by reference,

"Synchronization of Recurring Records in Incompatible

5 Databases", Serial no. 08/752,490, filed on November 13, 1996 (hereinafter, "application '490"); "Synchronization of Databases with Record Sanitizing and Intelligent Comparison," Serial no. 08/749,926, filed November 13, 1996 (hereinafter, "application '926"); "Synchronization of  
10 Databases with Date Range," Serial no. 08/748,645, filed November 13, 1996 (hereinafter, "application '645").

In the described embodiment, the modules of L\_translator 5 are designed specifically for interacting with local database 13 and local application 17. The design  
15 of the modules of L\_translator 5 is specifically based on the record and field structures and the rules of data value imposed on them by the local application, the Application Program Interface (API) requirements and limitations of local application 17 and other characteristics of the local  
20 database and application. The same is true of the modules of R\_translator 9. These translators are typically not able to interact with other databases or Applications and are only aware of the characteristics of the database and application for which they are designed. Therefore, when  
25 the user chooses two applications for synchronization, Translation Engine 1 chooses the two translators which are able to interact with those applications. In an alternate embodiment, the translators can be designed as table-driven codes, where a general translator is able to interact with a  
30 variety of applications and databases based on supplied parameters.

Having described the structure of synchronization program 100 in reference to its various modules, we will now

describe the operation of synchronization program 100.  
During synchronizing the two database, Control Module 2  
instructs the various modules in synchronization program 100  
to perform specific tasks. We will describe the operation  
5 of synchronization program 100 by describing the steps taken  
by Control Module 2 (as set out in the pseudo code in Fig.  
4) and describing in detail the actions by the various  
modules as they are instructed by Control Module 2.

Referring to Fig. 4, in the first step of  
10 synchronizing the two databases, Control Module 2 instructs  
the Parameter Table Generator 3 to create parameter table 4  
(Step 100). In this step, as part of creating parameter  
table 4, Parameter Table Generator 3 obtains from the user a  
filter expression, if any, to be used during synchronization  
15 or alternatively accesses a previously stored filter  
expression, if any.

Fig. 5 is pseudo code for the steps taken by  
Parameter Table Generator 3 in relation to filter  
expressions. In step 150, Parameter Table Generator 3  
20 determines from the user the whether a filter should be used  
during the current synchronization. Fig. 6 shows a filter  
selection window 40 which the user uses to select whether to  
use a filter for the current synchronization. If the user  
selects to use a filter (step 151), Parameter Table  
25 Generator 3 next determines the filter to be used (step  
152). This filter may be a previously stored filter  
expression (e.g. filter 42 in Fig. 6) or a filter expression  
which the user inputs (e.g. by selecting "New" button 44 in  
Fig. 6).

30 For a new filter, the user uses a filter name input  
window 50 to input the filter's name (step 153), shown in  
Fig. 7. The user then uses a filter criteria input window  
60 (step 154), shown in Figs. 8 and 9, to input the filter

expression for the new filter. We will describe window 60 in further detail below. Parameter Table Generator 3 next stores the current date and time in the FILTER\_CHANGED\_TIMESTAMP parameter (step 155). This parameter is used to determine whether a filter has changed since a previous synchronization. Parameter Table Generator 3 then assigns a unique filter identifier code to the filter expression, which is then used to identify the filter (step 156).

10           If the user selects to use a previously stored filter expression, Parameter Table Generator 3 obtains the name of the filter from the user (step 157) and then retrieves the filter expression and the unique filter identifier code of that filter (step 158). If the user  
15 selects to edit the filter, Parameter Table Generator 3 displays the filter expression in window 60 and allows the user to edit the filter expression (step 159). In this step, Parameter Table Generator 3 also stores the current date and time in the FILTER\_CHANGED\_TIMESTAMP parameter, as  
20 in step 155.

          In step 160, Parameter Table Generator 3 sets the FILTER\_ID parameter to the unique filter identifier of the filter that was selected. The modules of synchronization program 100 use FILTER\_ID parameter to determine the correct  
25 filter to use. Parameter Table Generator 3 next sets USE\_FILTER flag (step 161). This flag indicates to various module of synchronization program 100 that a filter is to be used during synchronization and causes the appropriate modules to take the necessary steps, as will be described in  
30 further detail below.

          During synchronization, the filter may be applied to the records of the database by either the translator for that database or by the database manager application of that

database, if the database manager application is capable of applying a filter. In step 162, Parameter Table Generator 3 parses the filter expression into a filter token array which the translators use when filtering records of the databases and history file. The filter token array and the manner in which it is used will be described in detail below.

Parameter table generator 3 will next create Parameter\_Table 4, as described in detail in the '490, '926 and '645 applications.

Having described the steps taken by Parameter Table Generator 3 with respect to a filter to be applied during synchronization, we will now describe in detail graphical user interface (GUI) windows displayed for entering the filter expressions, the semantics of filter language used in the described embodiment, the manner in which inputted filter criteria are stored, and the method used by translators to determine whether a record passes the filter. However, it should be noted that other filtering languages and methods may also be used to filter records during synchronization.

Generally, synchronization program 100 uses two types of filters: static and dynamic filters. Static filters have a fixed and unchanging filter expression. For example, the filter "appointments in 1997" is a static filter. It will always cover the appointments in 1997. Dynamic filters have a changing filter threshold value, which may depend on a changing parameter. For example, the filter "appointments from today until a year from today" is a dynamic filter because the threshold value of the filter changes as the value of "today" changes. (It should be noted that the above examples are also examples of date range filters. Date range filters filter records based on



whether the dates of the records fall within a range of dates specified by the filter.)

In the case of dynamic filters, Parameter Table Generator 3 uses the dynamic filter to create two filter expressions to be used during synchronization. The first filter expression, which we will refer to as the current filter, is based on the value of the filter for the current synchronization. For example, in the case of date range filters based on the value of the current synchronization's date, the value of the current filter will be based on the current synchronization's date. (Date range filters and a method of synchronizing databases using them are described in detail in the '490, '926 and '645 applications.) Since a dynamic filter is a changing filter, records which were previously within the filter may not be within the current filter. However, those records and any changes in those records should be used during the current synchronization since the user likely treated those records as being validly within the filter and might have modified them. Therefore, Parameter Table Generator 3 creates a second filter expression, which we will refer to as the loading filter, which combines the value the current filter with the value of the filter as it was during a previous synchronization. For example, in the case of a dynamic date range, the loading filter would be a concatenation of the current filter and the filter based on the date of the previous synchronization. The manner in which these two filters are used will be described below.

In synchronization program 100, a filter expression applied to both local and remote databases. However, the filter expression is typically inputted and stored based on the list of fields of one of the databases - in the described embodiments, the local database. A field map

which maps the fields of the two databases onto one another is used by synchronization program 100 to apply a filter expression based on the field list of one of the databases to the fields in the records of the other database, as will be described in detail below.

The table in Fig. 10 shows the specification of the semantics of the filter language. We define a filter expression (or filter criteria) as a group of one or more filter conditions (or filter criterion). In the filter language described here, when a record is evaluated against a filter condition, the result may be either a boolean value (i.e. TRUE or FALSE) or a numeric value (e.g. 1, 2, 3). However, the final result of evaluating a record against the filter expression is a boolean value which indicates whether the record has passed or failed the filter.

A filter condition may be considered to be a sentence having three parts: <evaluated operand> <filter operator> <filter threshold operand>, where the evaluated operand is the operand to be evaluated against the filter, filter threshold operand is the threshold value of the filter condition, and the operator is the filtering function to be performed between the evaluated operand and filter threshold operand. For example, the following is a filter condition: <date field> <is greater than or equal to> <TODAY>.

An operand may have one of two values: a value inputted by a user or a value taken from a record to be evaluated. To indicate that the value of an operand is to be taken from a field in a record, the name of that field is used as the operand. In the described filter language, a field name is enclosed with brackets - e.g. [Start Date]. During evaluation of a record against the filter, the value

stored in that field of the evaluated record will be used as the operand.

In the case where the operand has a user inputted value, the operand contains that value - e.g. 'Smith', 456-7896, or 'TODAY'. In the described filter language, the user-inputted operands may have one of the following types of values, which is coextensive with the possible field values of the local and remote databases: DATE, TIME, TEXTSTRING, BOOLEAN or NUMBER.

We will now describe in detail an example of the range of values the various types of user-inputted operands in the described filter language may have. However, it should be noted that other embodiments may have other ranges and limitations. In the described embodiment, DATE operands are formatted as 'YYYYMMDD' (in some embodiments, single quotes must be included) - example '19980101' is New Years Day of 1998. DATE operands may also have the value TODAY, which indicates the date for today obtained from the operating system of the computer. TIME operands are formatted as 'HHMM' on a 24-hour clock basis (in some embodiments, single quotes must be included) - for example, '0600' represents 6:00 AM and '1300' represents 1:00 PM. TIME operands may also have the value NOW, which indicates the current time obtained from the operating system of the local computer.

TEXTSTRING operands can contain any text value (in some embodiments, single quotes must be included to indicate that the value is a text string) - examples are 'Puma Technology, Inc.', '15', and '#\$%'. BOOLEAN operands must be of value TRUE or FALSE (no use of single quotes). NUMBER operands may be any integer or floating point number (in some embodiments, single quotes are not used for NUMBER operands).

In the described filter language, the available filter operators are organized into seven different operator sets. The first operator set (also referred to as "OP\_SET\_1") includes the following filter operators: EQ  
5 (equal), LE (less than or equal), GE (greater than or equal), NE (not equal), LT (less than) and GT (greater than). This set of operators may be used for all of the various types of operands, provided that both operands involved in the filter condition are of the same type.

10 The second operator set (also referred to as "OP\_SET\_2") is to be used only when evaluated operand is of the type DATE. OP\_SET\_2 provides for using dynamic filters for operands of type DATE (e.g. dynamic date range filters). OP\_SET\_2 is made up of all combinations of the filter  
15 operators in OP\_SET\_1 and the variables TODAY - and TODAY +. An OP\_SET\_2 filter operator is followed by a filter threshold operand whose value is a selected number of days. The variable TODAY + in an OP\_SET\_2 filter operator then indicates the number of days in the filter threshold operand  
20 is to be added to the date of the current synchronization prior using the OP\_SET\_1 filter operator to evaluate filter. For example, consider the filter condition: <appointment date> <LT TODAY +> <3>. In this filter expression, 3 days are added to the date of the current synchronization to  
25 obtain the date of the third day after today and then the filter operator LT is used to determine whether the appointment date is less than the date of the third day after today.

The third operator set (also referred to as  
30 "OP\_SET\_3") is similar the second operator set and is used only when both operands are of the type TIME. OP\_SET\_3 provides for using dynamic filters for operands of type TIME. OP\_SET\_3 includes all combinations of filter operators

in OP\_SET\_1 with the variables NOW - and NOW +. The variable NOW represents the time of the current synchronization, typically obtained from the operating system. An OP\_SET\_3 filter operator is followed by a filter  
5 threshold operand whose value is a selected number of seconds. The variables NOW + and NOW - are used in the same manner as the variables TODAY + and TODAY -.

The fourth operator set (also referred to as "OP\_SET\_4") may be used when both operands are of the type  
10 TEXTSTRING. OP\_SET\_4 includes the following operators: STARTS\_WITH, CONTAINS, DOES\_NOT\_CONTAIN, IS\_EMPTY and IS\_NOT\_EMPTY.

The fifth operator set (also referred to as "OP\_SET\_5") may be used when both operands are of the type  
15 NUMBER. OP\_SET\_5 includes the following operators: + (addition), - (subtraction), \* (multiplication), / (division) and % (modulus).

The sixth operator set (also referred to as "OP\_SET\_6") may be used only when both operands are of the  
20 type BOOLEAN. OP\_SET\_6 includes only the operator IS.

The seventh operator set (termed OP\_SET\_7) may be used to combine two filter conditions. OP\_AND\_OR includes the relational operators AND and OR. Filter conditions may be combined using the seventh operator set to form filter  
25 expressions. In the described embodiment, the operand from the seventh operator set are used so as to achieve one of two results: a record must either meet all of the filter conditions in a filter expression or only one of the filter conditions. In other embodiments, more complex filter  
30 expressions may be permitted. In such embodiments, the order of evaluation may follow a predetermined order of evaluation (e.g. the order of evaluation in the 'C'

programming language) which may in turn be modified by parentheses.

Following are several exemplary filter expressions based on the above described filter language:

- [Full Name] CONTAINS 'Smith' AND [Private Flag] IS FALSE
- ( [Priority] / 2 ) GT 0
- [Alarm Date] EQ '19980101' AND ( [Regarding] CONTAINS 'meeting' OR [Regarding] CONTAINS 'training' ) AND [Location] CONTAINS 'Boston'

Referring back to Figs. 8 and 9, the user uses the filter expression input window 60 to enter the filter expression to be used during synchronization. Filter expression input window 60 includes two so-called tabs (e.g. used in operating systems sold under the tradename Windows by Microsoft Corporation of Redmond, Washington). Conditions tab 62 (shown in Fig. 8) is used to input the filter conditions. Rules tab 64 (shown in Fig. 9) is used to select how the filter conditions should be combined to create the filter expression.

Referring to Fig. 8, conditions tab 62 includes a filter condition input region 66. The user can select a field name from a list of field names of the database on which the filter is based by using pull-down menu 66a. Alternatively, the user may type a valid name in field name region 66b. The user may also type in a valid filter threshold operand in threshold operand region 66e. In filter operator region 66c, the user may type a valid filter operator or select one from a filter operator pull-down menu (not shown; only button 66d for clicking on to pull-down the menu is shown). It should be noted that the list of available options in pull-down menu 66d is limited by the

type of operand entered in field name region 66b or filter threshold operand region 66e.

The user may add a filter condition to the list displayed in filter conditions display region 66g or remove  
5 a filter condition from that list, by selecting the appropriate button in region 66f. The user may also update (i.e. change or edit) a filter condition, by selecting the appropriate "button" in region 66f.

Referring to Fig. 9, in rules tab 64, more  
10 particularly in region 64a, the user may select whether a record must either meet all of the filter conditions inputted by the user or only one of the filter conditions in order to pass the filter.

The filter expression input by the user in the  
15 filter expression input window 60 is then stored as a filter expression based on the above described filter language.

We will now describe the filtering methodology used in synchronization program 100 to evaluate a record against a filter. Briefly, the filtering methodology used in  
20 synchronization program 100 generally has two steps. The first step is parsing the stored filter expression and forming a filter token array which is designed to facilitate evaluating the records against the filter. In synchronization program 100, Parameter Table Generator 3  
25 performs this step (Fig. 5, step 162). The second step is evaluating each record of the database or the history file to be filtered against the filter expression to determine whether the record passes the filter. In synchronization program 100, the translators performs this step in the case  
30 of the local and remote databases. Synchronizer 15 performs this step for the history file. We will now describe each of these steps in detail.

Parameter Table Generator 3 parses the stored filter expression and then forms the filter expression into a filter token array to be used during evaluation of the records against the filter expression. A token in the filter token array is a data structure which represents either an operand (i.e an evaluated operand or a filter threshold operand) or a filter operator. A token contains two pieces of information - the type of the token and the value of the token. The type of the token may be one of the following: TEXTSTRING, DATE, TIME, BOOLEAN, NUMBER, OP\_SET\_1, OP\_SET\_2, OP\_SET\_3, OP\_SET\_4, OP\_SET\_5, OP\_SET\_6 or OP\_SET\_7. The value of the token will be the actual content of the token and is stored as a string of characters which is obtained from the filter expression. For example, a token which represents a date filter threshold operand "2/7/1988" will have a DATE type and a "2/7/1988" value. A filter operand GE (greater than or equal to) will have a OP\_SET\_1 type and a "GE" value. An evaluated operand having a field name "[contact address]" will have a TEXTSTRING type and a "[contact address]" value.

As Parameter Table Generator 3 parses the filter expression, it turns the operands and operators of the filter conditions into tokens. Each filter condition yields three tokens. Parameter Table Generator 3 stores the tokens in the filter token array in a specific order. In the case of each filter condition, the evaluated operand token is stored first; followed by the filter threshold operand token and then the filter operator token. In the case of filter operators from OP\_SET\_7 (i.e., the filter operators AND and OR), the two filter conditions connected by the operator are stored first, followed by the filter operator. Ordering the tokens in this manner facilitates evaluating the records against the filter, as will be described below.



As an example, parameter generator 3 parses the following filter expression:

[Alarm Date] EQ '19980101' AND ( [Regarding]  
CONTAINS 'meeting' OR [Regarding] CONTAINS 'training' ) AND  
5 [Location] CONTAINS 'Boston'

into the following token array:

	<u>value</u>	<u>type</u>
	AND	OP_SET_7
	CONTAINS	OP_SET_4
10	'Boston'	TEXTSTRING
	[Location]	TEXTSTRING
	AND	OP_SET_7
	OR	OP_SET_7
	CONTAINS	OP_SET_4
15	'training'	TEXTSTRING
	[Regarding]	TEXTSTRING
	CONTAINS	OP_SET_4
	'meeting'	TEXTSTRING
	[Regarding]	TEXTSTRING
20	EQ	OP_SET_1
	'19980101'	DATE
	[Alarm Date]	DATE

As stated above, the second step in filtering the records of a database is evaluating each record of that  
25 database against the filter expression. The translator of the database to be filtered performs this step. Both of these modules use the same method of evaluation, which we will now describe.

To evaluate a record against a filter expression, a  
30 translator uses a stack, which we will refer to as the evaluation stack. To evaluate the record, starting with the last item in the filter token array, the translator proceeds through the filter token array and pushes copies of the

operand tokens onto the evaluation stack. When an operand token is a field name, the translator instead of pushing the field name, pushes the data stored in that field of the record onto the stack. When the translator encounters an operator token in the token array, the translator pops the last two items in the evaluation stack and evaluates the two items based on the filter operator. The translator push the result of the evaluation, which may be either of the type BOOLEAN or NUMBER, onto the evaluation stack. After the translator evaluates the entire filter expression, a single boolean value is left in the stack which indicates whether the record has passed the filter.

We will now provide an example of evaluating a record against a filter. Prior to the evaluating step, Parameter Table Generator 3 would have parsed the following filter expression:

[Last Name] STARTS\_WITH 'A' AND [City] EQ 'Boston'  
into the following token array:

<u>value</u>	<u>type</u>
AND	OP_SET_7
EQ	OP_SET_1
'Boston'	TEXTSTRING
[City]	TEXTSTRING
STARTS_WITH	OP_SET_4
'A'	TEXTSTRING
[Last Name]	TEXTSTRING

During filter evaluation, the first two tokens are pushed onto the evaluation stack:

<u>token array</u>	<u>evaluation stack</u>
AND	[Last Name]
EQ	'A'
'Boston'	

[City]

STARTS\_WITH

Next, the operator STARTS\_WITH is encountered. Therefore, two operands are popped from the evaluation stack and evaluated using the operator. The result of this evaluation (e.g. TRUE) is pushed back onto the evaluation stack:

<u>token array</u>	<u>evaluation stack</u>
AND	TRUE
EQ	
'Boston'	
[City]	

Two more operands from token array are pushed onto the evaluation stack:

<u>token array</u>	<u>evaluation stack</u>
AND	TRUE
EQ	[City]
	'Boston'

The operator EQ is next encountered. Therefore, two operands are popped from the evaluation stack and evaluated using the operator. The result of this evaluation (e.g. TRUE) is then pushed onto the evaluation stack.

<u>token array</u>	<u>evaluation stack</u>
AND	TRUE
	TRUE

Finally, the remaining operator AND is encountered in the filter token array. Two operands are popped from the evaluation stack and evaluated using the operator. The result of this evaluation (i.e. TRUE) is also pushed onto the evaluation stack. Because there are no more tokens in the filter token array, the remaining token on the evaluation stack is the final result of the filter

evaluation. This final token indicates that the record being evaluated passes the filter.

As stated briefly above, synchronization program 100 applies a filter expression based the field list of the remote database to the local database. To do so, synchronization program 100 uses a field map to determine which field of the local database corresponds to a field of the remote database used in the filter expression. Field mapping is described in U.S. Patent No. 5,392,390, incorporated by reference. Briefly, to synchronize records of two databases, it is essential to determine which field or fields of one database should be synchronized with which field or fields of the other database. To accomplish this, a field map is used which correlates the fields of the two databases to one another. It should be noted that not all fields of a database are mapped onto the other database and therefore such unmapped fields are not synchronized with the other database.

In the described embodiment, L\_translator 5 uses a filter expression which is based on the field list of the remote database to filter the records of the local database. For every token in the filter token array that contains a field name, L\_translator 5 uses the remote database to the local database field map to determine the corresponding field in the local database record being evaluated and pushes the content of that field onto the evaluation stack. It may be the case that the field in the filter expression is an unmapped field and therefore does not have counterpart in the local database record being evaluated. If that is the case, L\_translator 5 marks the token and its corresponding operator and operand tokens with a SKIP\_EVAL flag. During the evaluation phase if an operator token is marked with a SKIP\_EVAL flag, L\_translator 5 determines the

result of that operation to be TRUE. (If the operation to be skipped is to return a NUMBER type value, L\_translator 5 determines the result to be '0' or zero. L\_translator 5 then marks the resulting token, and that token's

5 corresponding operator token and other operand token, with a SKIP\_EVAL flag.) L\_translator 5 applies the filter to the record in this manner until a final result remains in the evaluation stack. In alternative embodiments, the filter expression may be based on the field list of the local  
10 database or the user may select the database on whose field list the filter expression is to be based.

As stated above, instead of the translator for the database to be filtered, the database manager application for that database may filter the records of that database.  
15 The translator for the database parses the filter expression into a set of instructions formatted for the Application Programmer Interface (API) of the database manager application. In that case, the database application manager transmits to synchronization program 100 only those records  
20 that pass the filter.

To determine whether the translator or the application will filter the records of a database, the modules of synchronization program 100 use two types of flags. First, as described above, Parameter Table Generator  
25 3 sets a USE\_FILTER flag if a filter is being used (Fig. 5, step 161). Second, each of translators 5 and 9 in turn set an appropriate flag, i.e. R\_Application\_Is\_Filtering or L\_Application\_Is\_Filtering, to indicate that the database manager application will apply the filter. If both the  
30 USE\_FILTER and the R\_Application\_Is\_Filtering (or L\_Application\_Is\_Filtering) flags are set, the flags indicate the remote (or local) database manager application will apply the filter to its database.

Referring back to Fig. 4, after Parameter Table Generator 3 creates the parameter table, Control Module 2 of the Translation Engine 1 instructs synchronizer 15 to initialize itself (step 101). Synchronizer 15 in response  
5 creates the workspace data array 16. Control Module 2 of the Translation Engine 1 then instructs synchronizer 15 to load history file 19 into workspace 16 (step 102). History file 19 is a file that was saved at the end of last synchronization and contains records reflecting the records  
10 of the two databases at the end of the previous synchronization. Synchronizer 15 uses history file 19 during current synchronization to analyze the records of the local and remote database to determine the changes, additions, and deletions in each of two databases since the  
15 previous synchronization. Synchronizer 15, as result of this analysis, then can determine what additions, deletions, or updates need be made to synchronize the records of the two databases.

In various situations, synchronizer 15 does not load  
20 history file 19. For example, if no history file from a previous synchronization exists or if the user chooses to synchronize not using the history file, synchronizer 15 will not load history file 19. Additionally, the user may wish to use a filter expression that is different from the filter  
25 expression used in the previous synchronization (including using no filter expression during the current synchronization). In that case, if one of the database manager applications filtered its database, then synchronizer 15 does not load the stored history file. This  
30 is because when a database manager application filters the records of a database, the history file contains only those records which pass the previous filter and not necessarily the records necessary for performing a synchronization using

the current filter. Obviously, in the case where a history file is not loaded, synchronizer 15 synchronizes the two databases without using a history file.

Fig. 11 is pseudocode for the steps taken by synchronizer 15 to load history file 19. For each Record in history file 19 (step 200), synchronizer 15 first loads the record (step 201) and then writes the loaded record into workspace 16 (step 202). Synchronizer 15 repeats these steps until all of the records of the history file are loaded into the workspace.

Referring back to Fig. 4, after the history file is loaded into the workspace, Control Module 2 instructs R\_translator 13 to load the remote database records (step 103). Fig. 12 is pseudocode for the steps taken by R\_translator 13 to load the remote database records. If the USE\_FILTER flag is set but R\_Application\_Is\_Filtering flag is not set (step 300) then R\_reader module 11 will apply the filter to the records of the remote database. For each record of the remote database (step 301), R\_reader module 11 of the R\_translator first loads the record (step 302). R\_reader module 11 applies the filter expression identified by the parameter Filter\_Id to the loaded record (step 303). If the record passes the filter then R\_reader module 11 marks the record as having passed the filter (step 304). If the record does not pass the filter then R\_reader module 11 marks the record as having failed the filter (step 305). R\_reader module 11 then sends the record to synchronizer 15 (step 306) and synchronizer 15 writes the loaded record into the workspace (step 307). Steps 302-307 are performed until all records of the remote database are loaded.

If the Use\_Filter flag and R\_Application\_Is\_Filtering flags are both set (step 309), then the remote database application will filter the loaded

records. R\_reader module 11 converts the filter expression into the format required by remote database application's API and sends the converted filter expression to the remote database application (step 310). R\_reader module 11 then  
5 loads the filtered records (step 311). If a record that was loaded passes the current filter expression, then R\_reader module 11 marks that as having passed the filter; otherwise, R\_reader module marks the record as having failed the filter (step 312). Since only those records that have passed the  
10 filter are loaded, R\_reader module 11 does not mark any of the loaded records as having failed the filter. R\_reader module 11 then sends the loaded records to synchronizer 15 (step 313) and synchronizer 15 writes the loaded records into the workspace (step 314).

15 Following loading the remote database records, Control Module 2 instructs L\_sanitizer module 8 of L\_translator 5 to sanitize the remote database records in the workspace (step 104).

Control Module 2 of the Translation Engine 1 then  
20 instructs the L\_translator 5 to load the records from the local database (step 105). L\_translator 5 and synchronizer 15 load records of the local database in the same manner as described for R\_translator 9 in reference to Fig. 12, except for two differences. First, as described above,  
25 L\_translator 5 filters the records of the local databases using the remote database to local database field map. It should be noted that the field maps are contained within the parameter table and the contents of the parameter table are transmitted to L\_translator as read-only data. Second, as  
30 synchronizer 15 receives each local database record from the L\_reader module 7 of the L\_translator 5, synchronizer 15 maps that record using the local database to remote database map before writing the record into the next available spot



in the workspace. This is due to the fact that, in the described embodiment, records in the workspace are stored according to the remote database data structure.

Referring back to Fig. 4, after all records are  
5 loaded into the workspace, Control Module 2 instructs synchronizer 15 to perform a Conflict Analysis and Resolution ("CAAR") procedure on the records in the workspace, which procedure is described in detail in the '490, '926 and '645 applications (step 107). Briefly,  
10 referring to Fig. 13, synchronizer 15 processes the records in the workspace, including comparing them to one another, in order to form them into groups of related records called corresponding item groups (CIGs). Synchronizer 15 forms the CIGs as it loads the records into the workspace and  
15 completes the process as the first step in CAAR (step 350). Each CIG may include at most one record from each of the databases and the history file. Each record in a CIG may be a recurring or a nonrecurring records. Where a database does not support recurring records, in the case of a  
20 recurring event, a CIG would contain related nonrecurring records from that database which together represent that recurring event. Based on this group of nonrecurring records, synchronizer 15 may create a model recurring record for use during the synchronization and include that model  
25 recurring record in the CIG. Hereinafter, when referring to a "record" in a CIG, we also intend to refer to such a group of related nonrecurring records in the CIG.

For each CIG (step 351), synchronizer 15 then compares the records in the CIG to one another, determines  
30 their differences, and decides what synchronization action should be taken (step 352). In essence, synchronizer 15 determines which record in the CIG contains the most current data. Synchronizer 15 then determine what synchronization

action should be taken to conform the other records in the CIG to the record with the most current data (i.e. how the other records in the CIG should be changed).

Synchronization actions with respect to a record include  
5 updating, deleting, adding, or not modifying that record.

We will now provide some examples of the results obtained in the CAAR analysis. If after comparing the records in a CIG, synchronizer 15 determines that the record from the local database is unchanged and the one from remote  
10 database is changed, synchronizer 15 determines that the local database record should be changed to conform to the remote database record. Or, if both records are changed (an example of what we refer to as a "conflict" since there is no clear choice of synchronization action), synchronizer 15  
15 may use a user-selected rule to decide what synchronization should be taken. The rule may require, for example, not modifying either of the records, changing the remote database record to conform to the local database record, or asking the user to resolve conflicts.

20 In the described embodiment, when a filter expression is used during the synchronization, synchronizer 15 alters the synchronization outcome in at least three cases. First, if the synchronization outcome is that a conflict exists (step 354), synchronizer 15 determines  
25 whether one of the records fails the current filter. If one of the records in the CIG fails the filter, then synchronizer 15 marks all records in the CIG so that they are not updated (step 355). In other embodiments, synchronizer may use the user-selected rule for resolving  
30 conflicts to resolve the conflict. If one of the records in the CIG does not fail the filter, then synchronizer 15 uses the user-selected rule to resolve the conflict (step 356).

Second, as stated above, synchronizer 15 changes the content of all records in a CIG to that of the record with the most up to date data. Therefore, if the record that contains the most up to date data fails the current filter, then the other records in the CIG when updated will also fail the filter. To resolve this issue, in the described embodiment, instead of updating the records in the CIG, all records in that CIG are marked as having failed filter (step 358). Therefore, the records are not updated when unloaded to the databases at the end of synchronization, as will be described below.

Third, the filter expression may contain a filter condition based on an unmapped remote database field. As described above, when applying such a filter to the local database records, the local translator evaluates the filter condition containing the unmapped field as having a TRUE value. In that case, some filtered local database records may be marked as having passed the filter while their corresponding remote database records may be marked as having failed the filter, even though the mapped fields of both records may contain the same data. In the described embodiment, if a such a filter expression is used and one of the records in a CIG is marked as having failed the current filter, then synchronizer 15 marks all of the CIG records as failing the filter (step 359).

As stated above, in some cases, one of the databases may support recurring records while the other database may not. Therefore, a recurring record is fanned into a number of non-recurring records before being unloaded to the database that does not support recurring records. In such a situation, during CAAR, synchronizer 15 examines each recurring record to determine whether there are some fanned nonrecurring records which pass the value of the dynamic

filter expression during the previous synchronization but fail the current value of the filter. If so, then the dynamic filter has changed in such a way that part of the set of fanned records fall outside of the current filter.

5 In the described embodiment, in such a situation, synchronizer 15 determines that the recurring record should be fanned again to generate new fanned nonrecurring records and previously fanned nonrecurring records should be deleted. To accomplish this, synchronizer 15 flags the  
10 recurring record and the appropriate translator fans the recurring record into the appropriate database and deletes the previous instances (step 360).

When synchronizer 15 finishes performing CAAR on the records, synchronizer 15 would have determined what  
15 synchronization action should be taken with respect all records to be synchronized. The records may then be unloaded into their respective databases. The translators will perform the specific synchronization actions to be taken with respect to the records of the databases.  
20 However, prior to doing so, the user is asked to confirm proceeding with unloading (Fig. 4, steps 108-109). Up to this point, neither the databases nor the history file have been modified. The user may obtain through the Translation Engine's User Interface various information regarding what  
25 synchronization actions will be taken upon unloading.

If the user chooses to proceed with synchronization and to unload, the records are then unloaded. The unloader modules 6,10 of the translators 5,9 perform the unloading for the databases. During unloading, translators may use  
30 the filter expression to limit the data that is unloaded to the databases. For example, the translators may unload only those records which fall within the filter expression and delete any record which falls outside of the filter

expression. During unloading, synchronizer 15 also creates the history file and unloads the records into the history file. We will now describe the unloading of the records into the databases and the history file in detail.

5           Control Module 2 of Translation Engine 1 first instructs R\_translator 9 to unload remote database records from workspace into the remote database (Fig. 4, step 110). Fig. 14 is pseudocode for the steps taken by R\_translator 9 to unload the records. For each remote database record in  
10 the workspace (step 400), R\_translator 9 first determines whether the Use\_Filter is set and the filter is a static filter (step 401). If that is the case, R\_translator 9 determines whether the record passes or fails the filter. (It should be noted that, as described above, some records  
15 are marked as failing the filter during CAAR. In that case, during unloading, those records are considered to fail the filter.)

          If the record fails the filter, R\_translator 9 deletes the record from the remote database. If the record  
20 is passes the filter, then R\_translator 9 adds, deletes, or modifies the record according to results of synchronization obtained during CAAR analysis (step 404). If the remote database does not support recurring records or in other circumstances described in detail in the '490, '926 and '645  
25 applications, R\_translator 9 in step 404 may fan a recurring record by creating an appropriate number of nonrecurring records corresponding to the recurring record. If, as described above, synchronizer 15 during CAAR marks a recurring record for re-fanning (Fig. 13, step 360),  
30 R\_translator 9 in this step will re-fan the record. When fanning, the number of nonrecurring records would be limited by any date based filters (i.e any date range) or other filters, so that nonrecurring records falling outside the

filter are not created. Additionally, if a user has selected to limit the number of fanned nonrecurring records for each recurring record, R\_translator would create only a limited number of instances, as described in more detail in the '490, '926 and '645 applications.

If the Use\_Filter is not set or the filter is not a static filter (step 401), R\_translator 9 determines whether the Use-Filter flag is set and the filter is a dynamic filter (step 405). If the Use-Filter flag is set and the filter is a dynamic filter, then R\_translator 9 determines whether the record to be unloaded passes or fails the current filter. If the record does not pass the current filter (step 406) then the record is deleted on the remote database (step 407). If the record passes the current filter, then R\_translator 9, in the same manner as step 404, adds, deletes, or modifies the record according to results of synchronization obtained during CAAR analysis (step 408).

By deleting records which fail the filter expression in steps 403 and 407, R\_translator 9 uses the filter to limit the size of the remote database. If the remote database is located on a handheld computer, R\_translator manages the memory of the handheld device by limiting the size of the database stored on the handheld computer.

Following unloading of the remote database records, Control Module 2 instructs the L\_translator to unload the local database records from the workspace (Fig. 4, step 111). Fig. 15 is pseudocode for the steps taken by L\_translator 5 to unload the local database records in the workspace. The steps 450-452, 454-455, 458-460 are respectively the same as steps 400-402, 404-405, and 408-410 in Fig. 14, described in detail above. Unlike R\_translator 9, L\_translator 5 does not delete records falling outside of the filter. Therefore, in step 453, if the filter is a

static filter and the record does not fit the filter, then L\_translator 5 modifies (i.e updates) the record if that is the synchronization result obtained during the CAAR analysis. However, synchronizer 15 does not add or delete a record, if that is synchronization result obtained during the CAAR analysis. In step 456, similarly, if the record is within the loading filter, then L\_translator 5 modifies the record if that is the synchronization result obtained during the CAAR analysis (step 457). In this manner, L\_translator 5 propagates to the records of the local database changes to those remote database records which do not fit the current filter expression. It should be noted that such remote database records are deleted from the remote database (step 407, Fig. 14) since those remote database records fail the current filter. Also, during unloading, the unloader module of the L\_translator uses the remote database to local database map to map the records in the workspace back into the format of local database records.

Additionally, it should be noted that where the local database manager application filters the local database, the local database manager application only provides those records which pass the loading filter. Therefore, the effect of step 457 is to update those records which passed the previous value of the filter expression but fail the current value of the filter expression.

It should be noted that in other embodiments, translators for the local and remote databases may use the filter expression during unloading in different manner than in the above described embodiments. For example, the remote translator may be configured in a similar manner as the local translator described above. Or, either the remote or local translator may only update records within the filter and leave completely unaffected records outside the filter.

A translator may also add new records from one database to another, if those records fall outside of the current filter but are within the loading filter.

Control Module 2 next instructs synchronizer 15 to  
5 create a new history file (step 112). The process of  
creating a history file is described in detail in the '490,  
'926 and '645 applications. Briefly, for each CIG,  
synchronizer 15 during the CAAR process determines which one  
of the records in the CIG should be saved as the history  
10 file record. Based on these results, synchronizer 15  
creates a history file. Synchronizer 15 also stores with  
each history file record the PASSED\_FILTER/FAILED\_FILTER  
flag based on whether the record passes or fails the current  
filter. Synchronizer also stores the value which determined  
15 the value of the dynamic filter for the current  
synchronization (e.g. the date of the current  
synchronization in the case of a dynamic date range filter).

At this point Synchronization is complete.

Other embodiments are within the following claims.

20 For example, if one of the databases has the  
capability to provide database generated information or data  
which can be used to determine, for example, whether a  
record has been changed, added, or deleted since a previous  
synchronization, the synchronization program uses that  
25 information to determine whether a record has been changed,  
added, or deleted. Of course, that database generated  
information is less than the whole record of the database.  
For example, that information may be a date and time stamp,  
or a flag, set when the record was last modified or when the  
30 record was added, whichever is later. We will briefly  
describe an embodiment of such a synchronization program,  
which is described in detail in the commonly assigned  
copending U.S. patent application, incorporated by reference



in its entirety, entitled "Synchronization of Databases,"  
filed on November 5, 1997, serial no. 08/964,751  
(hereinafter, the "'751 application").

There are generally two types of such databases:  
5 "medium synchronization" and "fast synchronization"  
databases. A "fast synchronization" database is a database  
which provides information regarding changes, deletions, and  
additions to its records from one synchronization to the  
next. A fast synchronization database also assigns to each  
10 record of the database a unique identification code (i.e. a  
unique ID) which uniquely identifies that record. Unique  
IDs are required to accurately identify records over a  
period of time. A fast synchronization database also  
provides a mechanism for keeping track of which records are  
15 added, changed, or deleted from synchronization to  
synchronization, including a list of deleted records.

A "medium synchronization" database typically has  
more limited capabilities than a fast synchronization  
database for keeping track of addition, deletions, or  
20 changes. In short, a medium synchronization database does  
not keep track of deletions. Such a database however still  
has the capability to provide information regarding what  
records were added or modified since a previous  
synchronization. A medium synchronization database also  
25 provides unique IDs.

If the information provided by a database indicates  
that a record has not been changed or added since a previous  
synchronization, the synchronization program need not load  
that record and can use the history file to reconstruct the  
30 relevant contents of that record for synchronizing the two  
databases. The history file contains a copy of the relevant  
content of that record as the record was at the time of  
(e.g. at the end of) the previous synchronization. Using

the history file to reconstruct the record instead of loading the record can result in significant saving of time - where for example the data transfer link between the two computers is slow - since typically a majority of records in  
5 databases are unchanged records. The synchronization program thereby increases the efficiency of performing synchronization between two databases.

The synchronization program does not synchronize a record of the fast or medium synchronization database that  
10 fails the filter expression with the records of the other database. Therefore, the synchronization program does not reconstruct those unchanged records which fail the filter expression. However, the synchronization reconstructs unchanged records of the fast or medium synchronization  
15 database which failed the filter during the last synchronization but pass the current filter. In that case, since the records previously failed the filter, the records would not be in the other database. After reconstructing these unchanged records, the synchronization program treats  
20 these records as if these records were newly added to the fast or medium synchronization database and therefore adds these record to the other database.

As is apparent, when synchronizing a fast or medium synchronization database, the synchronization program may  
25 use the history file to reconstruct unchanged records, whether those unchanged records fail or pass the current filter. Therefore, the synchronization program at the end of each synchronization ensures that the records of the history file are synchronized with the records of the fast  
30 or medium synchronization database. In essence, the synchronization program ensures that the history file contains records which represent all of the records of the fast or medium synchronization at the end of the current

synchronization, whether the records pass or fail the current filter. Since, as described above, some records of the fast or medium synchronization database are not present in the other database, the history file contains some  
5 records that are present only in the fast or medium synchronization database but not in the other database.

Where the database manager application of a fast or medium synchronization database filters the records of the database, the synchronization program does not receive those  
10 records of the database which fail the filter. Therefore, if the filter changes such that some of the unchanged records which were previously outside of the filter are now within the filter, the synchronization program can not rely on the history file. In that case, the synchronization  
15 program loads all records of the database and proceeds to synchronize without using the history file.

In some embodiments, both computers on which the two databases run are capable of running programs other than a database, as in the case of, for example, general purpose  
20 computers such as desktop and notebook computers, or handheld computers having sufficient memory and processing power. In such a case, the synchronization program may be distributed between the two computers so as to, for example, increase the efficiency of using of a slow data transfer  
25 link between the two machines. We will briefly describe an embodiment of such a distributed synchronization program, which is described in detail in the commonly assigned copending U.S. patent application, incorporated herein in its entirety by reference, entitled "Distributed  
30 Synchronization of Databases", filed on September 11, 1997, serial no. 08/927,922 (hereinafter, the "'922 application").

Briefly, referring to Fig. 16, at remote computer 22, remote segment 26 of the synchronization program loads

records of remote database 13. Remote segment 26 then determines which records of the remote database have been changed/added, deleted or left unchanged since a previous synchronization. If the remote database assigns unique  
5 identification codes (i.e. unique ID) to its records, remote segment 26 can further differentiate between records than have been added and those than have been changed since the previous synchronization. Remote segment 26 uses a remote history file 30 which stores data representing or reflecting  
10 the records of the database at the completion of the previous synchronization. This data may be a copy of remote database 13. It may also be hash numbers for each of the records of the remote database. If the remote database assigns unique IDs, the remote history file may contain  
15 those unique IDs together with the hash numbers of the records corresponding to the stored unique IDs.

Remote segment 26 sends those records of the remote database that have been changed or added to the host segment or the host computer. However, the remote segment does not  
20 send the unchanged or deleted records to the host computer. Instead, the remote segment sends a flag indicating the status of the record (e.g. unchanged or changed) and some data or information that uniquely identifies the record to the host segment. This data or information may be a hash  
25 number of all or selected fields in the record at the completion of the last synchronization. It may also be the unique ID assigned to the record by the remote database, if the database assigns one to its records.

Host segment 28 uses the received information or  
30 data that uniquely identifies the unchanged record to access a record in host history file 19 that corresponds to the received information or data. This record contains a copy of the data of the remote database record that the remote

segment found to have been unchanged. Host segment 19 then uses this record to synchronize the databases in the same manner as described above. After synchronization, the remote and host history files are updated. Also, the records are unloaded to the remote and local database based on the filter expression, in the same manner as described above. Since the unchanged records which typically constitute most of the records of a database are not transferred to the host computer, a data transfer link, specially a slow data transfer link, is used with increased efficiency.

In such a distributed synchronization program, the remote and host segments would apply the filter expression to the records of the databases. In the case of the host segment, the process would be similar to that for the above described embodiments. In the case of the remote database, along with information identifying the record or the record's field values, remote segment 26 sends the host segment information indicating that the record passed or failed the filter. The synchronization process then proceeds as for the previously described embodiments. During unloading, the host segment sends the remote segment information with respect whether the records passed or failed the filter expression, along with the result of CAAR. Remote segment 26 then uses this information and the filter expressions, in the same manner as the above described translators when unloading records to the remote database.

In some embodiments, two or more databases on one computer may be synchronized with one database on the other computer. For example, it may be that the remote database application supports only one name and address database while the local database application supports two name and address databases. To synchronize the multiple local

databases with the single remote database, one of the local databases is designated as the main local database for synchronization with the remote database. During synchronization, as synchronization program 100 adds records  
5 from the local databases to the remote database, synchronization program tags the added records with codes (i.e. origin tags) identifying the source of the database record, e.g. the first local database, the second local database, etc. Synchronization program 100 uses these tags  
10 during future synchronizations to ensure that the tagged records are synchronized with the correct database, i.e. the database from which they originated. Additionally, during synchronization, synchronization adds new records from the remote database only to the local database which was  
15 designated as the main database. This method of synchronization is described in detail the '490, '926 and '645 applications.

To enable filtering the records during synchronization, synchronization program 100 uses the origin tags of the records to ensure that the correct filter is  
20 applied to the correct records. For example, consider the case where two local databases are synchronized with a single remote database. Each of the local databases may have a unique filter. Or one database may have a filter and  
25 the other may not. When synchronization program 100 is synchronizing the local databases with the remote database, synchronization program 100 uses the origin tags of the remote database records to determine which filter should be applied to each record. If the origin tag of a remote  
30 database record indicates that it originated from the first database, then the filter expression for that database, if any, is used. Similarly, if the origin tag of a remote database record indicates that it originated from the second

database, then the filter expression for that database, if any, is used. Additionally, if a remote database record is new (i.e. newly added since the previous synchronization), the filter expression for the local database that was

5 designated as the main local database is used. It should be noted that while this method of synchronization was described for synchronizing a single remote database with multiple local database, same method may be used for synchronizing multiple remote databases with a single local  
10 database, or multiple remote databases with multiple local databases.

It should be noted that the synchronization process in the above embodiments was described primarily in reference to using a filter during synchronization. If a  
15 user chooses not to use a filter, the synchronization proceeds generally in the manner described the '751, '922, '490, '926 and '645 applications.

What is claimed is: